

Piecewise Constant Policy Approximations to Hamilton-Jacobi-Bellman Equations

C. Reisinger* and P.A. Forsyth†

January 21, 2016

Abstract

An advantageous feature of piecewise constant policy timestepping for Hamilton-Jacobi-Bellman (HJB) equations is that different linear approximation schemes, and indeed different meshes, can be used for the resulting linear equations for different control parameters. Standard convergence analysis suggests that monotone (i.e., linear) interpolation must be used to transfer data between meshes. Using the equivalence to a switching system and an adaptation of the usual arguments based on consistency, stability and monotonicity, we show that if limited, potentially higher order interpolation is used for the mesh transfer, convergence is guaranteed. We provide numerical tests for the mean-variance optimal investment problem and the uncertain volatility option pricing model, and compare the results to published test cases.

Key words: fully nonlinear PDEs, monotone approximation schemes, piecewise constant policy time stepping, viscosity solutions, uncertain volatility model, mean variance

AMS subject classification: 65M06, 65M12, 90C39, 49L25, 93E20

1 Introduction

This article is concerned with the numerical approximation of fully nonlinear second order partial differential equations of the form

$$0 = F(\mathbf{x}, V, DV, D^2V) = \begin{cases} V_\tau - \sup_{q \in Q} L_q V, & \mathbf{x} \in \mathbb{R}^d \times (0, T], \\ V(\mathbf{x}) - \mathcal{G}(\mathbf{x}), & \mathbf{x} \in \mathbb{R}^d \times \{0\}, \end{cases} \quad (1.1)$$

where $\mathbf{x} = (S, \tau)$ contains both ‘spatial’ coordinates $S \in \mathbb{R}^d$ and *backwards time* τ . For fixed q in a control set Q , L_q is the linear differential operator

$$L_q V = \text{tr}(\sigma_q \sigma_q^T D^2 V) + \mu_q^T DV - r_q V + f_q, \quad (1.2)$$

where $\sigma_q \in \mathbb{R}^{d \times d}$, $\mu_q \in \mathbb{R}^d$ and $r_q, f_q \in \mathbb{R}$ are functions of the control as well as possibly of \mathbf{x} . An initial (in backwards time) condition $V(0, \cdot) = \mathcal{G}(\cdot)$ is also specified.

*Mathematical Institute, Andrew Wiles Building, University of Oxford, Woodstock Road, Oxford, OX2 6GG, reisinge@maths.ox.ac.uk

†Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada N2L 3G1, paforsyt@uwaterloo.ca

These equations arise naturally from stochastic optimization problems. By *dynamic programming*, the value function satisfies an HJB equation of the form (1.1). Since dynamic programming works backwards in time from a terminal time T to today $t = 0$, it is conventional to write PDE (1.1) in terms of backwards time $\tau = T - t$, with T being the terminal time, and t being forward time.

Many examples of equations of the type (1.1) are found in financial mathematics, including the following: optimal investment problems [32]; transaction cost problems [17]; optimal trade execution problems [1]; values of American options [25]; models for financial derivatives under uncertain volatilities [2, 30]; utility indifference pricing of financial derivatives [15]. More recently, enhanced oversight of the financial system has resulted in reporting requirements which include Credit Value Adjustment (CVA) and Funding Value Adjustment (FVA), which lead to nonlinear control problems of the form (1.1) [12, 31, 13].

If the solution has sufficient regularity, specifically for Cordes coefficients, it has recently been demonstrated that higher order discontinuous Galerkin solutions are possible [36]. Generally, however, these problems have solutions only in the viscosity sense of [16].

A general framework for the convergence analysis of discretization schemes for strongly nonlinear degenerate elliptic equations of type (1.1) is introduced in [7], and has since been refined to give error bounds and convergence orders, see, e.g., [4, 5, 6]. The key requirements that ensure convergence are consistency, stability and monotonicity of the discretization.

The standard approach to solve (1.1) by finite difference schemes is to “discretize, then optimize”, i.e., to discretize the derivatives in (1.2) and to solve the resulting finite-dimensional control problem. The nonlinear discretized equations are then often solved using variants of policy iteration [20], also known as Howard’s algorithm and equivalent to Newton’s iteration under common conditions [9].

At each step of policy iteration, it is necessary to find the globally optimal policy (control) at each computational node. The PDE coefficients may be sufficiently complicated functions of the control variable q such that the global optimum cannot be found either analytically or by standard optimization algorithms. Then, often the only way to guarantee convergence of the algorithm is to discretize the admissible control set and determine the optimal control at each node by exhaustive search, i.e., Q is approximated by finite subset $Q_H = \{q_1, \dots, q_J\} \subset Q$. This step is the most computationally time intensive part of the entire algorithm. Convergence to the exact solution is obtained by refining Q_H .

Of course, in many practical problems, the admissible set is known to be of *bang-bang* type, i.e., the optimal controls are a finite subset of the admissible set. Then the true admissible set is already a discrete set of the form Q_H .

In both cases, if we use backward Euler timestepping, an approximation to V^{n+1} at time τ^{n+1} is obtained from

$$\frac{V^{n+1} - V^n}{\Delta\tau} - \max_{q_j \in Q_H} L_{q_j}^h V^{n+1} = 0, \quad (1.3)$$

where we have a spatial discretization $L_{q_j}^h$, with h a mesh size and $\Delta\tau$ the timestep.

1.1 Objectives

It is our experience that many industrial practitioners find it difficult and time consuming to implement a solution of equation (1.3). As pointed out in [34], many plausible discretization

schemes for HJB equations can generate incorrect solutions. Ensuring that the discrete equations are monotone, especially if accurate *central differencing as much as possible* schemes are used, is non-trivial [38]. Policy iteration is known to converge when the underlying discretization operator for a fixed control is monotone (i.e., an M-matrix) [9]. Seemingly innocent approximations may violate the M-matrix condition, and cause the policy iteration to fail to converge.

A convergent iterative scheme for a finite element approximation with quasi-optimal convergence rate to the solution of a strictly elliptic switching system is proposed and analysed in [10]. Here, we are concerned with parabolic equations and exploit the fact that approximations of the continuous-time control processes by those piecewise constant in time and attaining only a discrete set of values, lead to accurate approximations of the value function.

A technique which seems to be not commonly used (at least in the finance community) is based on piecewise constant policy time stepping (PCPT) [28, 6]. In this method, given a discrete control set $Q_H = \{q_1, \dots, q_J\}$, J independent PDEs are solved at each timestep. Each of the J PDEs has a constant control q_j . At the end of the timestep, the maximum value at each computational node is determined, and this value is the initial value for all J PDEs at the next step.

Convergence of an approximation in the timestep has been analyzed in [27] using purely probabilistic techniques, which shows that under mild regularity assumptions a convergence order of $1/6$ in the timestep can be proven. In this and other works [26, 28], applications to fully discrete schemes are given and their convergence is deduced. These estimates seem somewhat pessimistic, in that we typically observe (experimentally) first order convergence.

Note that this technique has the following advantages:

- No policy iteration is required.
- Each of the J PDEs has a constant policy, and hence it is straightforward to guarantee a monotone, unconditionally stable discretization.
- Since the PCPT algorithm reduces the solution of a nonlinear HJB equation to the solution of a sequence of linear PDEs (at each timestep), followed by a simple max or min operation, it is straightforward to extend existing (linear) PDE software to handle the HJB case.
- Each of the J PDEs can be advanced in time independently. Hence this algorithm is an ideal candidate for efficient parallel implementation.
- In the case where we seek the solution of a Hamilton-Jacobi-Bellman-Isaacs (HJBI) PDE of the form

$$V_\tau - \inf_{p \in P} \sup_{q \in Q} L_{q,p} V = 0, \quad (1.4)$$

the *discretize and optimize* approach may fail due to the fact that policy iteration may not converge in this case [37]. However, the PCPT technique can be easily applied to these problems.

In view of the advantages of piecewise policy time stepping, it is natural to consider some generalizations of the basic algorithm. Since each of the J independent PDE solves has a different control parameter, it is clearly advantageous to use a different mesh for each PDE solution. This may involve an interpolation operation between the meshes. If we restrict attention to purely monotone schemes, then only linear interpolation can be used.

However, in [6], it is noted that the solution of the PDE (1.3) can be approximated by the solution of a switching system of PDEs with a small switching cost. There, it is shown that the solution of the switching system converges to the solution of (1.3) as the switching cost tends to zero. In [6], the switching system was used as a theoretical tool to obtain error estimates.

Building on the work in [6], the main results of this paper are:

- We formulate the PCPT algorithm in terms of the equivalent switching system, in contrast to [27]. We then show that a non-monotone interpolation operation between the switching system meshes is convergent to the viscosity solution of (1.3). The only requirement is that the interpolation operator be of *positive coefficient* type. This permits use of limited high order interpolation or monotonicity preserving (but not monotone) schemes.
- We will include two numerical examples. The first example is an uncertain volatility problem [30, 2] with a bang-bang control, where we demonstrate the effectiveness of a higher order (not monotone) interpolation scheme. The second example is a continuous time mean-variance asset allocation problem [39]. In this case, it is difficult to determine the optimal policy at each node using analytic methods, and we follow the usual program of discretizing the control and determining the optimal value by exhaustive search. We compare the numerical results obtained using PCPT and a standard policy iteration method. Comparable accuracy is obtained for both techniques, with the PCPT method having a considerably smaller computational complexity.

1.2 Outline

In order to aid the reader, we provide here an overview of the steps we will follow to carry out our analysis. We will write the PCPT algorithm in the unconventional form

$$\frac{V_j^{n+1} - \max_{k=1}^J V_k^n}{\Delta\tau} - L_{q_j}^h V_j^{n+1} = 0, \quad (1.5)$$

where the optimization step is carried out at the beginning of the new timestep, as opposed to the conventional form whereby the optimization is performed at the end of the old timestep. Note that the scheme is a time-implicit discretization for each fixed control q_j , while the optimization is carried out explicitly. As discussed above, a decided advantage of this approach is that this scheme is unconditionally stable and yet no nonlinear iterations are needed in every timestep.

In order to carry out our analysis, we perform the following string of approximations:

$$\boxed{\text{HJB equation}} \quad V_\tau - \sup_{q \in Q} L_q V = 0 \quad (1.6)$$

$$\boxed{\text{Control discretization}} \quad V_\tau - \max_{q_j \in Q_H} L_{q_j} V = 0 \quad (1.7)$$

$$\boxed{\text{Switching system}} \quad \min(V_{j,\tau} - L_{q_j} V_j, V_j - \max_{k \neq j} (V_k - c)) = 0 \quad (1.8)$$

$$\boxed{\text{Discretization}} \quad \min \left(\frac{V_j^{n+1} - V_j^n}{\Delta\tau} - L_{q_j}^h V_j^{n+1}, V_j^{n+1} - \max_{k \neq j} (\tilde{V}_{k,(j)}^{n+1} - c) \right) = 0 \quad (1.9)$$

In the HJB equation (1.6), the control parameter q is assumed to take values in a compact set Q , and for fixed q , L_q is a second order elliptic operator as per (1.2). The compact set is discretized by a finite set $Q_H = \{q_1, \dots, q_J\}$, where H is the maximum distance between any element in Q and Q_H . Of course, in the case of a *bang-bang* control, the admissible set is already discrete.

The resulting equation (1.7) can then be approximated by the *switching system* (1.8) as in [6] when the cost $c > 0$ of switching between controls $j = 1, \dots, J$ goes to zero. When $c \rightarrow 0$, every V_j converges to the solution of (1.7). The switching cost is included to guarantee that the system (1.8) satisfies the no-loop condition [24], and hence a comparison property holds. We then freeze the policies over time intervals of length $\Delta\tau$, i.e., restrict the allowable policies to those that assume one of the q_j over such time intervals, and discretize the PDEs in space and time. Here, we use the same timestep $\Delta\tau$ for the, say, backward Euler time discretization of the PDE, but generalizations are straightforward. We provide for the possibility that the PDEs for different controls are solved on different meshes, and in that case interpolation of the discretized value function V_k for control q_k onto the j -th mesh is needed. We denote by $\tilde{V}_{k,(j)}$ this interpolant of V_k evaluated on the j -th mesh.

The remainder of this article is organized as follows. We conclude this section by giving standard definitions and assumptions on the equation (1.1). Section 2 shows that the control space can be approximated by a finite set, which prepares the formulation as a switching system. Section 3 introduces a discretization based on piecewise constant policy timestepping, while Section 4 contains the main result proving convergence of these approximation schemes satisfying a standard set of conditions, to the viscosity solution of a switching system. Section 5 constructs numerical examples for the mean-variance asset allocation problem and the uncertain volatility option pricing model. Section 6 concludes.

1.3 Preliminaries

We now give the standard definition of a viscosity solution before making assumptions on F . Given a function $f : \Omega \rightarrow \mathbb{R}$, where $\Omega \subset \mathbb{R}^n$ open, we first define the upper semi-continuous envelope as

$$f^*(\mathbf{x}) = \lim_{r \rightarrow 0^+} \sup \{ f(\mathbf{y}) \mid \mathbf{y} \in B(\mathbf{x}, r) \cap \Omega \}, \quad (1.10)$$

where $B(\mathbf{x}, r) = \{ \mathbf{y} \in \mathbb{R}^n \mid |\mathbf{x} - \mathbf{y}| < r \}$. We also have the obvious definition for a lower semi-continuous envelope $f_*(\mathbf{x})$.

Definition 1 (Viscosity Solution). *A locally bounded function $U : \Omega \rightarrow \mathbb{R}$ is a viscosity subsolution (respectively supersolution) of (1.1) if and only if for all smooth test functions $\phi \in C^\infty$, and for all maximum (respectively minimum) points \mathbf{x} of $U^* - \phi$ (respectively $U_* - \phi$), one has*

$$\begin{aligned} & F_*(\mathbf{x}, U^*(\mathbf{x}), D\phi(\mathbf{x}), D^2\phi(\mathbf{x}), U^*(\mathbf{x})) \leq 0 \\ & \left(\text{respectively} \quad F^*(\mathbf{x}, U_*(\mathbf{x}), D\phi(\mathbf{x}), D^2\phi(\mathbf{x}), U_*(\mathbf{x})) \geq 0 \right). \end{aligned} \quad (1.11)$$

A locally bounded function U is a viscosity solution if it is both a viscosity subsolution and a viscosity supersolution.

Remark 1 (Smoothness of test functions). *Definition 1 specifies that $\phi \in C^\infty$, whereas the common definition uses $\phi \in C^2$. The equivalence of these two definitions is discussed in [3, 35]. Letting $\phi \in C^\infty$ simplifies the consistency analysis.*

Assumption 1 (Properties of $F(\cdot)$). *We assume that Q is a compact set and that $\sigma_q \sigma_q^T, \mu_q, r_q, f_q$ are bounded on $\mathbb{R}^{d+1} \times Q$, Lipschitz in \mathbf{x} uniformly in q (i.e., there is a Lipschitz constant which holds for all q) and continuous in q .*

Remark 2 (Comparison principle). *Assumption 1 is the same as the one made in [6]. It guarantees (see, e.g., [6, 19]) that a strong comparison principle holds for F , i.e., if V and W are viscosity sub- and supersolutions, respectively, of (1.1), with $V(\cdot, 0) \leq W(\cdot, 0)$, then $V \leq W$ everywhere. It also ensures the well-posedness of the switching system (1.8), see [24] and [6].*

2 Approximation by finite control set

In this section, we analyze the validity of the first approximation step from (1.6) to (1.7), i.e., that the compact control set may be approximated by a finite set. More precisely, for compact $Q \subset \mathbb{R}^m$ (i.e., $m \in \mathbb{N}$ is the dimension of the parameter space) and $Q_h \subset Q$ such that

$$\max_{q \in Q} \min_{\hat{q} \in Q_h} |q - \hat{q}| \leq h, \quad (2.1)$$

we define a discrete control HJB equation by

$$0 = F_h(\mathbf{x}, V, DV, D^2V) = \begin{cases} V_\tau - \sup_{q \in Q_h} L_q V, & \mathbf{x} \in \mathbb{R}^d \times (0, T], \\ V(\mathbf{x}) - \mathcal{G}(\mathbf{x}), & \mathbf{x} \in \mathbb{R}^d \times \{0\}. \end{cases} \quad (2.2)$$

The following lemma will be useful.

Lemma 1 (Properties of $F(\cdot)$). *Under Assumption 1, for any \mathbf{x} and any C^∞ test function ϕ*

$$\begin{aligned} |F(\mathbf{x}, \phi(\mathbf{x}), D\phi(\mathbf{x}), D^2\phi(\mathbf{x})) - F_h(\mathbf{x}, \phi(\mathbf{x}) + \xi, D\phi(\mathbf{x}), D^2\phi(\mathbf{x}))| &\leq \omega_h(\mathbf{x}, h) + \omega_\xi(\xi), \\ \text{where } \omega_h(\mathbf{x}, h) &\rightarrow 0 \text{ as } h \rightarrow 0, \\ \omega_\xi(\xi) &\rightarrow 0 \text{ as } \xi \rightarrow 0, \end{aligned} \quad (2.3)$$

where $\omega_h(\mathbf{x}, h)$ is locally Lipschitz continuous in \mathbf{x} , independent of h .

Proof. See Appendix A. □

Theorem 1. *Given Assumption 1, let V and V_h be the unique viscosity solutions to*

$$F(\mathbf{x}, V, DV, D^2V) = 0, \quad (2.4)$$

$$F_h(\mathbf{x}, V_h, DV_h, D^2V_h) = 0. \quad (2.5)$$

Then $V_h \rightarrow V$ uniformly on compact sets as $h \rightarrow 0$.

Proof. A consequence of Lemma 1 is that

$$F_h^*(\mathbf{x}, \phi(\mathbf{x}) + \xi, \dots) \leq F^*(\mathbf{x}, \phi(\mathbf{x}), \dots) + \omega_h(\mathbf{x}, h) + \omega_\xi(\xi). \quad (2.6)$$

Define, for all \mathbf{x} ,

$$\underline{V}_h(\mathbf{x}) = \liminf_{\mathbf{y} \rightarrow \mathbf{x}} V_h(\mathbf{y}), \quad (2.7)$$

$$\underline{V}(\mathbf{x}) = \liminf_{\substack{h \rightarrow 0 \\ \mathbf{y} \rightarrow \mathbf{x}}} \underline{V}_h(\mathbf{y}). \quad (2.8)$$

We claim that \underline{V} is a viscosity supersolution of equation (2.4). To show this, fix $\hat{\mathbf{x}}$ and choose a smooth test function ϕ such that $\hat{\mathbf{x}}$ is a global minimum of $\underline{V} - \phi$, and that

$$\underline{V}(\hat{\mathbf{x}}) = \phi(\hat{\mathbf{x}}). \quad (2.9)$$

Then, there exists a sequence $\hat{\mathbf{x}}_h \rightarrow \hat{\mathbf{x}}$, $h \rightarrow 0$, $\underline{V}_h(\hat{\mathbf{x}}_h) \rightarrow \underline{V}(\hat{\mathbf{x}})$, such that $\hat{\mathbf{x}}_h$ is a global minimum of $\underline{V}_h(\hat{\mathbf{x}}_h) - \phi(\hat{\mathbf{x}}_h)$. At each point $\hat{\mathbf{x}}_h$, since \underline{V}_h is a viscosity solution of (2.5),

$$0 \leq F_h^*(\hat{\mathbf{x}}_h, \underline{V}_h(\hat{\mathbf{x}}_h), D\phi(\hat{\mathbf{x}}_h), D^2\phi(\hat{\mathbf{x}}_h)). \quad (2.10)$$

Let $\underline{V}_h(\hat{\mathbf{x}}_h) = \phi(\hat{\mathbf{x}}_h) + \xi_h$, $\xi_h \rightarrow 0$, $h \rightarrow 0$, so that equation (2.10) becomes

$$0 \leq F_h^*(\hat{\mathbf{x}}_h, \phi(\hat{\mathbf{x}}_h) + \xi_h, D\phi(\hat{\mathbf{x}}_h), D^2\phi(\hat{\mathbf{x}}_h)). \quad (2.11)$$

From equations (2.6) and (2.11) we obtain

$$0 \leq F^*(\hat{\mathbf{x}}_h, \phi(\hat{\mathbf{x}}_h), D\phi(\hat{\mathbf{x}}_h), D^2\phi(\hat{\mathbf{x}}_h)) + \omega_h(\mathbf{x}, h) + \omega_\xi(\xi_h), \quad (2.12)$$

which gives us

$$\begin{aligned} 0 &\leq \limsup_{\substack{h \rightarrow 0 \\ \hat{\mathbf{x}}_h \rightarrow \hat{\mathbf{x}}}} F^*(\hat{\mathbf{x}}_h, \phi(\hat{\mathbf{x}}_h), D\phi(\hat{\mathbf{x}}_h), D^2\phi(\hat{\mathbf{x}}_h)) + \limsup_{\substack{h \rightarrow 0 ; \xi_h \rightarrow 0 \\ \mathbf{x} \rightarrow \hat{\mathbf{x}}}} (\omega_h(\mathbf{x}, h) + \omega_\xi(\xi_h)) \\ &\leq F^*(\hat{\mathbf{x}}, \phi(\hat{\mathbf{x}}), D\phi(\hat{\mathbf{x}}), D^2\phi(\hat{\mathbf{x}})) \\ &= F^*(\hat{\mathbf{x}}, \underline{V}(\hat{\mathbf{x}}), D\phi(\hat{\mathbf{x}}), D^2\phi(\hat{\mathbf{x}})), \end{aligned} \quad (2.13)$$

where we use equation (2.9) and Lemma 1.

Using similar steps, we can show that \bar{V} defined similar to (2.8) is a viscosity subsolution of equation (2.4). Invoking the strong comparison principle, $\bar{V} = \underline{V} = V$. Uniform convergence on compact sets follows using the same argument as in Remark 6.4, of [16]. \square

3 Piecewise constant policy timestepping

Here, we use the equivalence of (1.7) and (1.8) established in [6] as $c \rightarrow 0$ to formulate (1.9) precisely. Consider the HJB equation

$$\begin{aligned} V_\tau &= \max_{q_j \in Q} L_{q_j} V, \\ Q &= \{q_1, q_2, \dots, q_J\}, \end{aligned} \quad (3.1)$$

where we assume a discrete set of controls Q . We have shown in Section 2 that the optimal value under controls chosen from a compact set can be approximated by a control problem with a finite set.

According to [6], we can also approximate equation (3.1) by a *switching system*. Let $U_j, j = 1, \dots, J$, be the solution of a system of HJB equations, with

$$\begin{aligned} \min \left[U_{j,\tau} - L_{q_j} U_j, U_j - \left(\max_{k \neq j} (U_k - c) \right) \right] &= 0, \quad \mathbf{x} \in \mathbb{R}^d \times (0, T], \\ U_j - \mathcal{G}(\mathbf{x}) &= 0, \quad \mathbf{x} \in \mathbb{R}^d \times \{0\}. \end{aligned} \quad (3.2)$$

The constant $c > 0$ is required in order to add some small transaction cost to switching from $j \rightarrow k$. This cost term also ensures that only a finite number of switches can occur (otherwise there would be an infinite transaction cost; see also Remark 8). It is shown in [6] that $U_j \rightarrow V$ as $c \downarrow 0$, for all j .

For computational purposes, we define a finite computational domain $\Omega \subset \mathbb{R}^d$. Let $\partial\hat{\Omega}$ denote the portions of $\partial\Omega$ where we apply approximate Dirichlet conditions. We use the usual notation for representing (3.2). Define

$$\mathbf{x} = (S, \tau), \quad DU = (U_\tau, U_S), \quad D^2U = U_{SS}, \quad (3.3)$$

and let $\mathcal{B}_j(\mathbf{x})$ be the approximate Dirichlet boundary conditions on $\partial\hat{\Omega}$. Then, the *localized* problem is defined as

$$\begin{aligned} 0 &= F_j(\mathbf{x}, U_j, D U_j, D^2 U_j, \{U_k\}_{k \neq j}) \\ &= \begin{cases} \min \left[U_{j,\tau} - L_{q_j} U_j, U_j - (\max_{k \neq j} (U_k - c)) \right], & \mathbf{x} \in \Omega \setminus \partial\hat{\Omega} \times (0, T], \\ U_j(\mathbf{x}) - \mathcal{G}(\mathbf{x}), & \mathbf{x} \in \Omega \times \{0\}, \\ U_j(\mathbf{x}) - \mathcal{B}_j(\mathbf{x}), & \mathbf{x} \in \partial\hat{\Omega} \times (0, T], \end{cases} \end{aligned} \quad (3.4)$$

for $j = 1, \dots, J$. Letting $p_j = D U_j$, $s_j = D^2 U_j$, we can write equation (3.4) as

$$F_j(\mathbf{x}, U_j, p_j, s_j, \{U_k\}_{k \neq j}) = 0. \quad (3.5)$$

Note that system (3.4) is quasi-monotone (see [23]), since

$$F_j(\mathbf{x}, U_j, p_j, s_j, \{U_k\}_{k \neq j}) \leq F_j(\mathbf{x}, U_j, p_j, s_j, \{W_k\}_{k \neq j}) \quad \text{if } U_k \geq W_k; \quad k \neq j. \quad (3.6)$$

We include here the definition of a viscosity solution for systems of PDEs of the form (3.5) as used in [23, 11, 24].

Definition 2 (Viscosity solution of switching system). *A locally bounded function $U : \Omega \rightarrow \mathbb{R}^J$ is a viscosity subsolution (respectively supersolution) of (3.5) if and only if for all smooth test functions $\phi_j \in C^\infty$, and for all maximum (respectively minimum) points \mathbf{x} of $U_j^* - \phi_j$ (respectively $U_{j*} - \phi_j$), one has*

$$\begin{aligned} &F_{j*}(\mathbf{x}, U_j^*(\mathbf{x}), D \phi_j(\mathbf{x}), D^2 \phi_j(\mathbf{x}), \{U_k^*(\mathbf{x})\}_{k \neq j}) \leq 0 \\ &\left(\text{respectively } F_j^*(\mathbf{x}, U_{j*}(\mathbf{x}), D \phi_j(\mathbf{x}), D^2 \phi_j(\mathbf{x}), \{U_{k*}(\mathbf{x})\}_{k \neq j}) \geq 0 \right). \end{aligned} \quad (3.7)$$

A locally bounded function U is a viscosity solution if it is both a viscosity subsolution and a viscosity supersolution.

Remark 3. *Note that the j -th test function only replaces the derivatives operating on U_j . The terms which are a function of $U_k, k \neq j$, are not affected.*

We discretize (3.4) using the idea of piecewise constant policy timestepping. Define a set of nodes $S_{j,i}$ and timesteps τ^n , with discretization parameters h and $\Delta\tau$, i.e.,

$$\begin{aligned} \max_{\substack{1 \leq j \leq J \\ S \in \Omega}} \min_i |S - S_{j,i}| &= h, \\ \max_n (\tau^{n+1} - \tau^n) &= \Delta\tau. \end{aligned} \quad (3.8)$$

The distinction between $\Delta\tau$ and h is useful for the formulation of the algorithm, but somewhat arbitrary for the analysis. We will therefore label meshes and approximations by h and assume that

$$\Delta\tau \rightarrow 0 \quad \text{as} \quad h \rightarrow 0.$$

Define

$$\mathbf{x}_{j,i}^n(h) = (S_{j,i}, \tau^n; h) \in \Omega_{j,h}, \quad (3.9)$$

where $(S_{j,i}, \tau^n)$ refer to points on a specific grid j , and the set of grid points on the grid parameterized by h is $\Omega_{j,h}$.

Then we denote the discrete approximation to $U_j(\mathbf{x}_{j,i}^n)$ on a grid parameterized by h by $u_j(h, \mathbf{x}_{j,i}^n)$, which is extended to a function $u_j(h, \cdot)$ on $\Omega \times \{\tau_n\}$ by interpolation. We will sometimes use the shorthand notation

$$u_{j,i}^n = u_j(h, \mathbf{x}_{j,i}^n), \quad \mathbf{x}_{j,i}^n = (S_{j,i}, \tau^n), \quad (3.10)$$

where the dependence on h is understood implicitly. Note that by parameterizing $\mathbf{x}_{j,i}^n$ by the control index j , we are allowing for different discretization grids for different controls.

Let $\mathcal{L}_{q_j}^h$ be the discrete form of the operator \mathcal{L}_{q_j} . We discretize equation (3.4) for $\mathbf{x} \in \Omega \setminus \partial\hat{\Omega} \times (0, T]$, using $\Delta\tau = \tau^{n+1} - \tau^n$ constant for simplicity, by *piecewise constant policy timestepping*

$$\begin{aligned} u_{j,i}^{n+\frac{1}{2}} &= \max \left[u_{j,i}^n, \max_{k \neq j} \left(\tilde{u}_{k,i(j)}^n - c \right) \right], \\ u_{j,i}^{n+1} - \Delta\tau L_{q_j}^h u_{j,i}^{n+1} &= u_{j,i}^{n+\frac{1}{2}}, \quad j = 1, \dots, J, \end{aligned} \quad (3.11)$$

where

$$\tilde{u}_{k,i(j)}^n \equiv u_k(h, \mathbf{x}_{j,i}^n) \quad (3.12)$$

is the value of this interpolant $u_k(h, \cdot)$ of $u_k(h, \mathbf{x}_{j,i}^n)$ at the i -th point of grid $\Omega_{j,h}$.

Discretization (3.11) applies the “max” constraint at the beginning of a new timestep. Conventionally, one thinks of piecewise constant policy timestepping as applying the constraint at the end of a timestep.

Clearly, these would be algebraically the same thing if $u_{j,i}^{n+\frac{1}{2}}$ instead of $u_{j,i}^n$ was considered as approximation to $U_j(\mathbf{x}_{j,i}^n)$. So at the final timestep, these two possible approximations only differ by a final max operation. However, it will be convenient to apply the constraint as in equation (3.11). We can then rearrange equation (3.11) to obtain an equation in the form

$$\begin{aligned} &G_j \left(\mathbf{x}_i^{n+1}, h, u_{j,i}^{n+1}, \{u_{j,a}^{b+1}\}_{\substack{a \neq i \\ \text{or } b \neq n}}, \{\tilde{u}_{k,i}^n\}_{k \neq j} \right) \\ &= \min \left[\frac{u_{j,i}^{n+1} - u_{j,i}^n}{\Delta\tau} - L_{q_j}^h u_{j,i}^{n+1}, u_{j,i}^{n+1} - \max_{k \neq j} \left(\tilde{u}_{k,i(j)}^n - c \right) - \Delta\tau L_{q_j}^h u_{j,i}^{n+1} \right] \\ &= 0, \quad \mathbf{x} \in \Omega \setminus \partial\hat{\Omega} \times (0, T]. \end{aligned} \quad (3.13)$$

In the event that L_q is strictly elliptic, then we can interpret the effect of switching at the beginning of the timestep as adding a vanishing viscosity term to the switching part of the equation. However, note that we do not in general require that L_q be strictly elliptic.

We omit the trivial discretizations of $F_j(\cdot)$ in the remaining (boundary) portions of the computational domain. Note that the notation

$$\{u_{j,a}^{b+1}\}_{\substack{a \neq i \\ \text{or } b \neq n}} \quad (3.14)$$

refers to the set of discrete solution values at nodes neighbouring (in time and space) node $(i, n+1)$.

4 Convergence of approximations to the switching system

Here, we prove convergence of the piecewise constant policy approximation (1.9) to the solution of the switching system (1.8). We start by summarizing the main conditions we need. We will verify that these conditions are satisfied for all the schemes we use in our numerical examples.

Condition 1 (Positive interpolation). *We require that the interpolant $\tilde{u}_{k,i(j)}^n$ of the k -th grid onto the i -th point of the j -th grid can be written as*

$$\tilde{u}_{k,i(j)}^n = \sum_{\alpha \in N^k(j,i,n)} \omega_{k,i(j),\alpha}^n u_{k,\alpha}^n = \sum_{\alpha \in N^k(j,i,n)} \omega_{k,i(j),\alpha}^n u_k(h, \mathbf{x}_{k,\alpha}^n), \quad (4.1)$$

where

$$\sum_{\alpha \in N^k(j,i,n)} \omega_{k,i(j),\alpha}^n = 1, \quad \omega_{k,i(j),\alpha}^n \geq 0, \quad (4.2)$$

and $N^k(j,i,n)$ are the neighbours to the point $\mathbf{x}_{j,i}^n$ on grid $\Omega_{k,h}$.

Remark 4 (Monotone versus limited interpolation). *Condition (4.1), (4.2) is clearly satisfied by linear interpolation on simplices and multi-linear interpolation on hyper-rectangles, in which case the weights $\omega_{k,i(j),\alpha}^n$ are only functions of the coordinates $x_{j,i}^n$ and $x_{k,\alpha}^n$. This interpolation is then also a monotone operation and results in an overall monotone scheme, as defined below in Condition 2.*

We can, however, relax the requirement of monotonicity of the interpolation step by allowing weights $\omega_{k,i(j),\alpha}^n = \omega_{k,i(j),\alpha}^n(u_k^n)$, i.e., possibly nonlinear functions of the interpolated nodal values. It is easy to see that as long as

$$\min_{\alpha \in N^k(j,i,n)} u_{k,\alpha}^n \leq \tilde{u}_{k,i(j)}^n \leq \max_{\alpha \in N^k(j,i,n)} u_{k,\alpha}^n, \quad (4.3)$$

one can then always find weights such that (4.1), (4.2) hold. One can enforce (4.3) by a simple limiting step applied to any, potentially higher order, interpolant.

Generally, the decomposition (4.1) will not be the way by which a particular interpolation is originally defined or constructed in practice. The point is that whenever (4.3) holds, weights of this form exist and this is all that is needed for the theoretical analysis. An example are one-dimensional monotonicity preserving interpolants such as those in [21]. As discussed in [18], the interpolation in [21] is constructed to be monotonicity preserving (i.e., the interpolant is increasing over intervals where the interpolated values are increasing), and therefore satisfies condition (4.2).

Condition 2 (Weak Monotonicity). We require that discretization (3.13) be monotone with respect to $u_{j,a}^{b+1}, \tilde{u}_{k,i}^n$, i.e., if

$$\begin{aligned} w_{j,i}^n &\geq u_{j,i}^n & \forall(i, j, n), \\ \tilde{w}_{k,i(j)}^n &\geq \tilde{u}_{k,i(j)}^n & \forall(i, k, n), \end{aligned} \quad (4.4)$$

then

$$G_j \left(\mathbf{x}_i^{n+1}, h, u_{j,i}^{n+1}, \{w_{j,a}^{b+1}\}_{\substack{a \neq i \\ \text{or } b \neq n}}, \{\tilde{w}_{k,i(j)}^n\}_{k \neq j} \right) \leq G_j \left(\mathbf{x}_i^{n+1}, h, u_{j,i}^{n+1}, \{u_{j,a}^{b+1}\}_{\substack{a \neq i \\ \text{or } b \neq n}}, \{\tilde{u}_{k,i(j)}^n\}_{k \neq j} \right) \quad (4.5)$$

Remark 5. Note that the requirement that the scheme be monotone in $\tilde{u}_{k,i}^n$ (the interpolated solution) is a weaker condition than requiring monotonicity in $u_{k,\alpha}^n$. In particular, let us re-iterate that we are not requiring interpolation to be a monotone operation, as long as Condition 1 is satisfied.

Condition 3 (l_∞ stability). We require that the solution of equation (3.13), $u_j(h, \mathbf{x}_{j,i}^{n+1})$, exists and is bounded independent of h .

Remark 6. The bounds (4.3) implied by Condition 1 ensure that the interpolation step does not increase the l_∞ norm of the solution.

Condition 4 (Consistency). We require that we have local consistency, in the sense that, for any smooth function ϕ_j , and any functions ρ_k (not necessarily smooth)

$$\begin{aligned} &\left| G_j \left(\mathbf{x}_{j,i}^{n+1}, h, \phi_{j,i}^{n+1} + \xi, \{\phi_{j,a}^{b+1}\}_{\substack{a \neq i \\ \text{or } b \neq n}} + \xi, \{\tilde{\rho}_k(\mathbf{x}_{j,i}^{n_\ell})\}_{k \neq j} \right) \right. \\ &\quad \left. - F_j \left(\mathbf{x}_{j,i}^{n+1}, \phi_j(\mathbf{x}_{j,i}^{n+1}), D \phi_j(\mathbf{x}_{j,i}^{n+1}), D^2 \phi_j(\mathbf{x}_{j,i}^{n+1}), \{\tilde{\rho}_k(\mathbf{x}_{j,i}^{n_\ell})\}_{k \neq j} \right) \right| \leq \omega_1(h) + \omega_2(\xi), \\ &\quad \omega_1(h) \rightarrow 0 \text{ as } h \rightarrow 0, \quad \omega_2(\xi) \rightarrow 0 \text{ as } \xi \rightarrow 0. \end{aligned} \quad (4.6)$$

Theorem 2. Under Assumption 1, the solution of any scheme of the form (3.13) satisfying Conditions 1–4 converges to the viscosity solution of (3.4), uniformly on bounded domains.

Proof. We basically follow along the lines in [7], with the generalizations in [11] for weakly coupled systems. We include the details here in order to show that we only require Condition 1 for the interpolation operator, and this permits use of certain classes of high order interpolation.

Define the upper semi-continuous function \bar{u} by

$$\bar{u}_j(\hat{\mathbf{x}}) = \limsup_{\substack{h \rightarrow 0 \\ \mathbf{x}_{j,i}^{n+1} \rightarrow \hat{\mathbf{x}}}} u_j(h, \mathbf{x}_{j,i}^{n+1}), \quad (4.7)$$

where $\mathbf{x}_{j,i}^{n+1} \in \Omega_{j,h}$. Similarly, we define the lower semi-continuous function \underline{u} by

$$\underline{u}_j(\hat{\mathbf{x}}) = \liminf_{\substack{h \rightarrow 0 \\ \mathbf{x}_{j,i}^{n+1} \rightarrow \hat{\mathbf{x}}}} u_j(h, \mathbf{x}_{j,i}^{n+1}). \quad (4.8)$$

Note that the above definitions imply that $\bar{u}_j^* = \bar{u}_j$ and $\underline{u}_{j*} = \underline{u}_j$.

Let $\hat{\mathbf{x}}$ be fixed and ϕ_j be a smooth test function such that

$$\begin{aligned}\phi_j(\hat{\mathbf{x}}) &= \bar{u}_j(\hat{\mathbf{x}}), \\ \phi_j(\mathbf{x}) &> \bar{u}_j(\mathbf{x}), \quad \mathbf{x} \neq \hat{\mathbf{x}}.\end{aligned}\tag{4.9}$$

This of course means that $(\bar{u}_j - \phi_j)$ has a global maximum at $\mathbf{x} = \hat{\mathbf{x}}$. Consider a sequence of grids with discretization parameter h_ℓ , such that $h_\ell \rightarrow 0$ for $\ell \rightarrow \infty$. We use the notation

$$\mathbf{x}_{j,i_\ell}^{n_\ell+1} = (S_{j,i_\ell}, \tau^{n_\ell+1}; h_\ell) \in \Omega_{j,h_\ell}\tag{4.10}$$

to refer to the grid point $(i_\ell, n_\ell + 1)$ on the grid parameterized by h_ℓ , associated with discrete control q_j . Let $\mathbf{x}_{j,i_\ell}^{n_\ell+1}$ be the point on grid Ω_{j,h_ℓ} such that

$$\left(u_j(h_\ell, \mathbf{x}_{j,i_\ell}^{n_\ell+1}) - \phi_j(\mathbf{x}_{j,i_\ell}^{n_\ell+1}) \right)\tag{4.11}$$

has a global maximum, where ϕ_j is a test function satisfying equation (4.9). Note that in general, for any finite h_ℓ , $\mathbf{x}_{j,i_\ell}^{n_\ell+1} \neq \hat{\mathbf{x}}$.

Following the usual arguments from [7], and more particularly in [11], for a sequence of grids Ω_{j,h_ℓ} parameterized by h_ℓ , there exists a set of grid nodes $(i_\ell, n_\ell + 1)$, such that (4.11) is a global maximum and

$$\mathbf{x}_{j,i_\ell}^{n_\ell+1} \rightarrow \hat{\mathbf{x}}, \quad \mathbf{x}_{j,i_\ell}^{n_\ell} \rightarrow \hat{\mathbf{x}}, \quad u_j(h_\ell, \mathbf{x}_{j,i_\ell}^{n_\ell+1}) \rightarrow \bar{u}_j(\hat{\mathbf{x}}), \quad \text{for } \ell \rightarrow \infty,$$

where $\mathbf{x}_{j,i_\ell}^{n_\ell+1}, \mathbf{x}_{j,i_\ell}^{n_\ell} \in \Omega_{j,h_\ell}$ and, for $k \neq j$, noting equation (4.2) for the interpolant $\tilde{u}_k(\mathbf{x}_{k,i_\ell(j)}^{n_\ell})$, we have

$$\begin{aligned}\tilde{u}_{k,i_\ell(j)}^{n_\ell} &\equiv \tilde{u}_k(h_\ell, \mathbf{x}_{j,i_\ell}^{n_\ell}) = \sum_{\alpha_\ell \in N^k(j, i_\ell, n_\ell)} \omega_{k,i_\ell(j), \alpha_\ell}^{n_\ell} u_{k, \alpha_\ell}^{n_\ell} = \sum_{\alpha_\ell \in N^k(\cdot)} \omega_{k,i_\ell, \alpha_\ell}^{n_\ell} u_k(\mathbf{x}_{k, \alpha_\ell}^{n_\ell}), \quad k \neq j, \\ &\sum_{\alpha_\ell \in N^k(\cdot)} \omega_{k,i_\ell, \alpha_\ell}^{n_\ell} \mathbf{x}_{k, \alpha_\ell}^{n_\ell} \rightarrow \hat{\mathbf{x}}, \quad \ell \rightarrow \infty, \\ &\limsup_{\ell \rightarrow \infty} \tilde{u}_{k,i_\ell(j)} \leq \bar{u}_k(\hat{\mathbf{x}}),\end{aligned}\tag{4.12}$$

where $\mathbf{x}_{k, \alpha_\ell}^{n_\ell} \in \Omega_{k,h_\ell}$ and $\hat{\mathbf{x}} = (\hat{S}, \hat{\tau})$. Let

$$\begin{aligned}u_{j,i_\ell}^{n_\ell+1} &\equiv u_j(h_\ell, \mathbf{x}_{j,i_\ell}^{n_\ell+1}), \\ \phi_{j,i_\ell}^{n_\ell+1} &\equiv \phi(\mathbf{x}_{j,i_\ell}^{n_\ell+1}).\end{aligned}\tag{4.13}$$

Note that for any finite mesh size h_ℓ the global maximum in equation (4.11) is not necessarily zero, hence we define ξ_ℓ by

$$u_{j,i_\ell}^{n_\ell+1} = \phi_{j,i_\ell}^{n_\ell+1} + \xi_\ell,$$

such that

$$\xi_\ell \rightarrow 0 \quad \text{for } \ell \rightarrow \infty.\tag{4.14}$$

Since (4.11) is a global maximum at $\mathbf{x}_{i_\ell}^{n_\ell+1}$, then

$$\{u_{j,a_\ell}^{b_\ell+1}\}_{\substack{a_\ell \neq i_\ell \\ \text{or } b_\ell \neq n_\ell}} \leq \{\phi_{j,a_\ell}^{b_\ell+1}\}_{\substack{a_\ell \neq i_\ell \\ \text{or } b_\ell \neq n_\ell}} + \xi_\ell. \quad (4.15)$$

Substituting equations (4.14) and (4.15) into equation (3.13), and using the monotonicity property of the discretization (4.4) gives

$$0 \geq G_j \left(\mathbf{x}_{j,i_\ell}^{n_\ell+1}, h_\ell, \phi_{j,i_\ell}^{n_\ell+1} + \xi_\ell, \{\phi_{j,a_\ell}^{b_\ell+1}\}_{\substack{a_\ell \neq i \\ \text{or } b \neq n}} + \xi_\ell, \{\tilde{u}_{k,i_\ell}^{n_\ell}\}_{k \neq j} \right). \quad (4.16)$$

Note that we do not replace $\{\tilde{u}_{k,i_\ell}^{n_\ell}\}_{k \neq j}$ by the test function in equation (4.16). This is because the test function is only defined such that $\phi_j \geq \bar{u}_j$, and there is no such relationship with $u_k, k \neq j$. Let

$$\tilde{\rho}_{k,i_\ell}^{n_\ell} = \tilde{\rho}_k(\mathbf{x}_{j,i_\ell}^{n_\ell}) = \max(\tilde{u}_k(\mathbf{x}_{j,i_\ell}^{n_\ell}), \bar{u}_k(\hat{\mathbf{x}})). \quad (4.17)$$

From equation (4.12), we have that

$$\lim_{\ell \rightarrow \infty} \{\tilde{\rho}_k(\mathbf{x}_{j,i_\ell}^{n_\ell})\}_{k \neq j} = \bar{u}_k(\hat{\mathbf{x}}), \quad (4.18)$$

and that

$$\tilde{\rho}_{k,i_\ell}^{n_\ell} \geq \tilde{u}_{k,i_\ell}^{n_\ell}. \quad (4.19)$$

Substituting equation (4.17) into equation (4.16), and using the monotonicity property of the discretization (4.4) gives

$$0 \geq G_j \left(\mathbf{x}_{j,i_\ell}^{n_\ell+1}, h_\ell, \phi_{j,i_\ell}^{n_\ell+1} + \xi_\ell, \{\phi_{j,a_\ell}^{b_\ell+1}\}_{\substack{a_\ell \neq i \\ \text{or } b \neq n}} + \xi_\ell, \{\tilde{\rho}_{k,i_\ell}^{n_\ell}\}_{k \neq j} \right). \quad (4.20)$$

Equations (4.20) and (4.6) then imply that

$$0 \geq F_j \left(\mathbf{x}_{j,i_\ell}^{n_\ell+1}, \phi_j(\mathbf{x}_{j,i_\ell}^{n_\ell+1}), D \phi_j(\mathbf{x}_{j,i_\ell}^{n_\ell+1}), D^2 \phi_j(\mathbf{x}_{j,i_\ell}^{n_\ell+1}), \{\tilde{\rho}_k(\mathbf{x}_{j,i_\ell}^{n_\ell})\}_{k \neq j} \right) - \omega_1(h_\ell) - \omega_2(\xi_\ell).$$

Recalling equation (4.18), we have that

$$\begin{aligned} 0 &\geq \liminf_{\ell \rightarrow \infty} F_j \left(\mathbf{x}_{i_\ell}^{n_\ell+1}, \phi_j(\mathbf{x}_{i_\ell}^{n_\ell+1}), D \phi_j(\mathbf{x}_{i_\ell}^{n_\ell+1}), D^2 \phi_j(\mathbf{x}_{i_\ell}^{n_\ell+1}), \{\tilde{\rho}_k(\mathbf{x}_{j,i_\ell}^{n_\ell})\}_{k \neq j} \right) \\ &\quad - \liminf_{\ell \rightarrow \infty} \omega_1(h_\ell) - \liminf_{\ell \rightarrow \infty} \omega_2(\xi_\ell) \\ &\geq F_{j*} \left(\hat{\mathbf{x}}, \phi_j(\hat{\mathbf{x}}), D \phi_j(\hat{\mathbf{x}}), D^2 \phi_j(\hat{\mathbf{x}}), \{\bar{u}_k(\hat{\mathbf{x}})\}_{k \neq j} \right) \\ &= F_{j*} \left(\hat{\mathbf{x}}, \bar{u}_j(\hat{\mathbf{x}}), D \phi_j(\hat{\mathbf{x}}), D^2 \phi_j(\hat{\mathbf{x}}), \{\bar{u}_k(\hat{\mathbf{x}})\}_{k \neq j} \right). \end{aligned} \quad (4.21)$$

Hence \bar{u}_j is a subsolution of equation (3.4). A similar argument shows that \underline{u}_j is a supersolution of equation (3.4). Since a strong comparison principle holds for the switching system under Assumption 1 (see Proposition 2.1 in [6] and Remark 2), we then have $\underline{u}_j = \bar{u}_j$ is the unique continuous viscosity solution of equation (3.4).

We have thus shown the result. \square

Remark 7. Note that we have $\{\tilde{u}_{k,i(j)}^n\}$ appearing in equation (3.13), which would appear to cause a problem in terms of consistency, since we cannot assume that $\tilde{u}_{k,i(j)}^n = \tilde{u}_{k,i(j)}^{n+1} + O(h)$ since $u_{k,\alpha}^n$ are not necessarily smooth. The key fact here is that since equation (4.18) holds, we do not need smoothness.

Remark 8 (Interpolation and switching cost). We include here a brief example of the role of the switching cost c in (3.2). Assume we solve the degenerate equation $u_t = 0$ and write it as the trivial HJB equation $\sup_{\theta \in \{0,1\}} u_t^\theta = 0$. We represent u^0 and u^1 on two different meshes with nodes at $x_i^\theta = (i + 0.5\theta)h$, i.e., shifted by half a mesh size. The fully implicit discretization over a single timestep is $u^{\theta,n+1} = u^{\theta,n}$ for both controls. Consider now the situation of convex (u_i^0) , such that linear interpolation increases the solution. Using piecewise linear interpolation with $c = 0$ at the end of the timestep,

$$u_i^{1,n+1} = (u_i^{0,n+1} + u_{i+1}^{0,n+1})/2,$$

and similarly for $u_i^{0,n+1}$. Repeating this for the next timestep, the piecewise constant policy discretization is equivalent to

$$u_i^{\theta,n+2} = \frac{1}{4}u_{i-1}^{\theta,n} + \frac{1}{2}u_i^{\theta,n} + \frac{1}{4}u_{i+1}^{\theta,n} \quad \Leftrightarrow \quad \frac{u_i^{\theta,n+2} - u_i^{\theta,n}}{(h/2)^2} = \frac{u_{i-1}^{\theta,n} - 2u_i^{\theta,n} + u_{i+1}^{\theta,n}}{h^2}.$$

If we pick $\Delta\tau = h^2/8$, this discretization is consistent with the standard heat equation $u_t = u_{xx}$ instead of $u_t = 0$.

With the cost $c > 0$ switched on, for sufficiently small h , we will have

$$u_i^{1,n} > (u_i^{0,n+1} + u_{i+1}^{0,n+1})/2 - c,$$

such that $u^{1,n+1} = u^{1,n}$ and, by induction, the equation $u_t = 0$ is solved exactly.

For an overall convergent method, one has to pick h and Δt as a function of c , depending on the interpolation method and smoothness of the solution, and then let $c \rightarrow 0$. We discuss this in more detail in Section 5 for concrete examples.

Remark 9 (Reduction to standard piecewise constant policy method [27, 6]). Discretization (3.11), with $c = 0$, can be viewed as a form of the usual piecewise constant policy method [27, 6]. As a result, if linear interpolation is used to transfer information between grids, then (3.11) is a monotone discretization of HJB equation (3.1), which is easily shown to satisfy the standard requirements for convergence to the viscosity solution. If we use standard finite difference schemes, we would expect the spatial error (for smooth test functions) to be of size $O(h)$ with timestepping error of size $O(\Delta\tau)$. In this case ($c = 0$), if the solution is smooth, we would expect the total discretization error to be of size $O(h) + O(\Delta\tau) + O(h^2/\Delta\tau)$, where the last term arises from a linear interpolation error accumulated $O(1/\Delta\tau)$ times. Note that this term will be absent in the case $c > 0$, since the finite switching cost will prevent an $O(1/\Delta\tau)$ accumulation of interpolation error.

5 Numerical examples

In this section, we study the convergence of discretization schemes based on piecewise constant policy timestepping in numerical experiments. We present two examples, the uncertain volatility model from derivative pricing, and a mean-variance asset allocation problem. In both examples, we investigate the convergence with respect to the timestep and mesh size.

r	σ_{\min}	σ_{\max}	T	K	K_1	K_2	S_0
0.05	0.3	0.5	1	100	80	120	100

TABLE 5.1: *Model parameters used in numerical experiments for uncertain volatility model.*

5.1 The uncertain volatility model

We study first the uncertain volatility option pricing model [30]. In this example, we also examine the role of the switching cost and the impact of different interpolation methods.

The super-replication value of a European-style derivative is given by the HJB equation

$$\frac{\partial V}{\partial \tau} - \sup_{\sigma \in \Sigma} L_{\sigma} V = 0, \quad (5.1)$$

where

$$L_{\sigma} V = \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + r S \frac{\partial V}{\partial S} - r V \quad (5.2)$$

for $S \in (0, \infty)$, $\tau \in (0, T]$, with $\tau = T - t$ and

$$\Sigma = [\sigma_{\min}, \sigma_{\max}]. \quad (5.3)$$

In addition to the PDE, the value function satisfies a terminal condition. For the numerical tests, we choose a butterfly payoff function P such that

$$V(S, 0) = P(S) = \max(S - K_1, 0) - 2 \max(S - K, 0) + \max(S - K_2, 0)$$

and we localize the domain to $[S_{\min}, S_{\max}]$ with

$$\begin{aligned} V(S_{\max}, \tau) &= 0, & (S, \tau) &\in \{S_{\max}\} \times (0, T], \\ V_{\tau} &= -rV, & (S, \tau) &\in \{S_{\min}\} \times (0, T], \end{aligned}$$

and parameters as in Table 5.1.

It is well understood [30] that the optimal control is always attained at one of the interval boundaries (‘bang-bang’), depending on the sign of the second derivative of the value function V . The payoff function was chosen such that V has mixed convexity and therefore the optimal control differs between different regions of the state space and changes in time.

We use logarithmic coordinates $X = \log S$, so instead of (5.2) we approximate

$$L_{\sigma} V = \frac{1}{2} \sigma^2 \frac{\partial^2 V}{\partial X^2} + (r - \sigma^2/2) \frac{\partial V}{\partial X} - rV, \quad (5.4)$$

and now the coefficients are bounded and $X \in [\log S_{\min}, \log S_{\max}]$, so Assumption 1 is satisfied. Here, it is straightforward to construct “positive coefficient” schemes for the linear PDE arising when the control set is a singleton. A well-established route allows us to construct monotone, consistent and ℓ_{∞} stable schemes for the fully non-linear problem from positive coefficient discretization operators L_q^h using a direct control method of the “discretize, then optimize” type, see [20],

$$\min_{j \in \{1, 2\}} \left(\frac{V^{n+1} - V^n}{\Delta \tau} - L_{q_j}^h V^{n+1} \right) = 0, \quad (5.5)$$

where $q_1 = \sigma_{\min}$ and $q_2 = \sigma_{\max}$.

The resulting non-linear finite dimensional system of equations can be solved by policy iteration (Howard's algorithm, see [9]). In particular, we will use standard central finite differences in space and an implicit Euler discretization in time. For small enough h , this scheme is monotone, ℓ_∞ stable and consistent in the viscosity sense.

We compare the direct control method to several variants of piecewise constant policy timestepping on the basis of these discretizations, as described in Section 3, specifically equation (3.11). Here, only two control values have to be considered and the switching system is two-dimensional. In the case of no interpolation, the scheme simplifies to

$$\frac{V_j^{n+1} - \max_{k \in \{1,2\}} (V_k^n - c_{k,j})}{\Delta t} = L_{q_j}^h V_j^{n+1}, \quad j = 1, 2, \quad (5.6)$$

$$c_{k,j} = \begin{cases} c, & k \neq j, \\ 0, & k = j. \end{cases}$$

We again discretize L_{q_j} using a positive coefficient discretization, hence it is straightforward to verify that Assumption 1 and Conditions 1–4 hold (on the localized domain $S \in [S_{\min}, S_{\max}]$).

A solution extrapolated from the finest meshes was computed as an approximation to the exact solution and used to estimate the errors. The numerical value of this solution is $V(S_0, 0) = 1.67012$ (see also [33]). From this, we approximate the error as

$$e(h, \Delta\tau) = |V(S_0, 0) - \tilde{U}_1(h, \Delta\tau, c; S_0, 0)|, \quad (5.7)$$

where V is the exact solution and $\tilde{U}_1(h, \Delta\tau, c; \cdot, \cdot)$ the numerical approximation to U_1 , the first component of the switching system, for mesh size h , timestep $\Delta\tau$, and switching cost c .

Dependence on timestep $\Delta\tau$ and switching cost c

We first analyze how the switching cost affects convergence of the approximations. In Fig. 5.1, we compare the following two cases.

1. **Policy timestepping, fixed mesh:** We use (5.6) on a single uniform mesh on $[\log(K) - 4 \cdot \bar{\sigma}, \log(K) + 4 \cdot \bar{\sigma}]$, i.e., encompassing four standard deviations either side, where $\bar{\sigma}$ is the average of the two extreme volatilities.

For fixed switching cost, the error first decreases linearly as the timestep decreases, but eventually converges to a non-zero value. This convergence is monotone. Moreover, for fixed (sufficiently small) h , the asymptotic error for $\Delta\tau \rightarrow 0$ is decreasing with c .

2. **Policy timestepping, linear interpolation:** We also study the use of separate meshes for the two components of the switching system, V_1 and V_2 . In particular, we use uniform meshes on the intervals $[\log(K) - 4 \cdot \sigma_{\min}, \log(K) + 4 \cdot \sigma_{\min}]$ and $[\log(K) - 4 \cdot \sigma_{\max}, \log(K) + 4 \cdot \sigma_{\max}]$, so that for the chosen parameters mesh points on the two meshes do not coincide. Then, interpolation is necessary to represent these solutions on both meshes and to evaluate the explicit terms (those at time-level $n + 1$) in (5.6). In the case of linear interpolation, the overall scheme is monotone, and trivially satisfies Condition 1. As a result of the switching cost, the cumulative effect of linear interpolation in each timestep is controlled even for small $\Delta\tau$ (see Remark 9).




FIGURE 5.1: The uncertain volatility test case with parameters as in Table 5.1. Shown is for different methods the \log_2 error, where $e_m = e((\Delta\tau)_m, h, c)$ from (5.7) is the error for timestep $(\Delta\tau)_m = 1/8 \cdot 2^{-m} \in \{1/8, \dots, 1/32768\}$ and in each plot, from top to bottom, $c = 1/20, 1/40, 1/80, 1/160, 1/320, 1/640$. The mesh size is fixed at $h = 1/1024$. The dotted line has slope -1 . Left: Piecewise constant policy timestepping on a single mesh; right: linear interpolation between individual meshes for each control.

Next, we analyze the convergence jointly in h and $\Delta\tau$ for fixed c , as well as the convergence with respect to c for the case with interpolation. The results are given in Table 5.2. For fixed positive switching cost $c > 0$, we compute approximations on a sequence of time and space meshes with $N_k = 2N_{k-1}$ and $M_k = 2M_{k-1}$. The asymptotic ratio of about two is consistent with an error of $O(h) + O(\Delta\tau)$.

We now study the difference between the solution of the switching system for fixed cost c and the solution of the HJB equation (5.1). Considering the boldface values in the table as good approximations for this difference, we observe convergence as $c \rightarrow 0$ which is roughly consistent with order $3/4$. The theoretically proven order of $1/3$ from Theorem 2.3 in [6] does not seem sharp for these data.

The approximation error in Δt and h does not appear to be affected by c as long as $c > 0$, which is seen by comparison of the lines ‘(c)’ in Table 5.2 for $c = 1/10$ to $c = 1/640$, for small enough mesh parameters (last three columns). In computations, it would therefore seem prudent to pick $\Delta t = O(h)$ and $c = O(h^{4/3})$ i.e., proportional errors, even though faster, but more erratic convergence is obtained setting $c = 0$ uniformly.

Dependence on timestep $\Delta\tau$ and mesh size h

In Fig. 5.2, we show the convergence in both the timestep and mesh size for two different costs, $c = 0.01$ and $c = 0.16$. Compare this to the case $c = 0$ in Fig. 5.3 (middle row, left). We can make a number of observations. For large switching cost (right plot), the difference between the switching system and the HJB equation is large and dominates the discretization error. For fixed c and h , there appears to be convergence as $\Delta\tau \rightarrow 0$, and if we also let $h \rightarrow 0$ the solutions converge to the solution of the switching system with fixed c . For comparable values for h and $\Delta\tau$, there

N_k		32	64	128	256	512	1024	2048	4096
M_k		512	1024	2048	4096	8192	16384	32768	65536
c									
1/10	(a)	2.0692	1.9724	1.9660	1.9532	1.9491	1.9478	1.9474	1.9472
	(b)	0.3991	0.3023	0.2958	0.2831	0.2789	0.2777	0.2773	0.2771
	(c)		-0.0968	-0.0065	-0.0127	-0.0042	-0.0012	-0.0004	-0.0002
	(d)			14.9669	0.5079	3.0621	3.4257	2.8297	2.3076
1/40	(a)	1.6126	1.5441	1.7406	1.7663	1.7623	1.7612	1.7608	1.7606
	(b)	-0.0575	-0.1260	0.0705	0.0961	0.0922	0.0910	0.0906	0.0905
	(c)		-0.0685	0.1965	0.0256	-0.0040	-0.0011	-0.0004	-0.0002
	(d)			-0.3486	7.6692	-6.4629	3.5076	2.7965	2.3037
1/160	(a)	1.1989	1.2205	1.5510	1.7019	1.7033	1.7022	1.7018	1.7016
	(b)	-0.4712	-0.4496	-0.1191	0.0318	0.0331	0.0320	0.0317	0.0315
	(c)		0.0216	0.3305	0.1509	0.0013	-0.0011	-0.0004	-0.0002
	(d)			0.0653	2.1902	112.7836	-1.2152	2.7969	2.2994
1/640	(a)	0.9256	0.9897	1.3752	1.6448	1.6833	1.6822	1.6818	1.6816
	(b)	-0.7446	-0.6804	-0.2949	-0.0254	0.0132	0.0121	0.0117	0.0115
	(c)		0.0641	0.3855	0.2696	0.0385	-0.0011	-0.0004	-0.0002
	(d)			0.1664	1.4301	6.9932	-35.1108	2.7835	2.3009
0	(a)	0.8271	0.7221	1.0227	1.4109	1.6654	1.6702	1.6703	1.6702
	(b)	-0.8430	-0.9480	-0.6474	-0.2592	-0.0047	0.0001	0.0002	0.0001
	(c)		-0.1050	0.3005	0.3882	0.2545	0.0048	0.0001	-0.0001
	(d)			-0.3494	0.7742	1.5252	53.3568	39.0446	-1.6580

TABLE 5.2: *The uncertain volatility test case with parameters as in Table 5.1; piecewise constant policy timestepping with linear interpolation between individual meshes for each control; convergence with respect to mesh parameters and switching cost. Shown are: (a) the numerical solution $V_k = \tilde{V}(N_k, M_k, c; S_0, 0)$; (b) the difference to the exact solution $V_k - V(S_0, 0)$; (c) the increments $V_k - V_{k-1}$; (d) the ratios of increments $(V_k - V_{k-1})/(V_{k-1} - V_{k-2})$.*




FIGURE 5.2: *The uncertain volatility test case with parameters as in Table 5.1. Piecewise constant policy timestepping with linear interpolation between individual meshes for each control. Shown is the \log_2 error, where $e_m = e((\Delta\tau)_m, h, c)$ from (5.7) is the error for timestep $(\Delta\tau)_m = 1/8 \cdot 2^{-m} \in \{1/8, \dots, 1/524288\}$ and in each plot, from top to bottom, $h = 1/4, 1/16, 1/64, 1/256, 1/1024, 1/4096$. The cost size is fixed at $c = 0.01$ (left) and $c = 0.16$ (right). The dotted line has slope -1 . The downward spikes are a result of error cancellation for a particular combination of h , $\Delta\tau$ and c .*

appears to be a cancellation of leading order errors in h and $\Delta\tau$ with opposite signs, which appears as downward spikes in the left and middle plot.

In this particular case of linear interpolation, since the overall scheme is monotone, convergence to the viscosity solution of (5.1) is ensured if $c = 0$ as long as the discretization is consistent. Recall from Remark 9, that (for smooth solutions) the discretization error is of the form $O(\Delta\tau) + O(h^2) + O(h^2/\Delta\tau)$, with the third term being the cumulative effect of linear interpolation over $O(1/\Delta\tau)$ timesteps. Hence we can ensure consistency by requiring that $\Delta\tau = O(h)$.

We now analyze in more detail the convergence in $\Delta\tau$ and h for the degenerate case with $c = 0$. The computational results are shown in Fig. 5.3 and are discussed in the following list.

1. **Policy timestepping, fixed mesh:** We use (5.6) on a single uniform mesh on $[\log(K) - 4 \cdot \bar{\sigma}, \log(K) + 4 \cdot \bar{\sigma}]$, i.e., encompassing four standard deviations either side, where $\bar{\sigma}$ is the average of the two extreme volatilities. For fixed mesh size, the error approaches a constant level for decreasing time-step, but for simultaneously diminishing mesh size the observed time discretization error is clearly of first order in the timestep.
2. **Direct control:** Here, the optimal control is implicitly found with the solution as described above – see, in particular, (5.5) – and therefore we can disentangle the Euler discretization error from the effect of piecewise constant control. Comparing the envelope to the curves, parallel to the dotted line with slope minus one, shows first order convergence as in the previous case, but with a lower intercept which indicates that the time discretization error is about a factor 4 smaller. Although the number of policy iterations per time-step was consistently small (usually 2–4), the computational time here was dramatically larger (due to the need to generate new matrices in each iteration) and therefore solutions could not be

computed for the same number of timesteps as for the other cases.

3. **Policy timestepping, linear interpolation:** Again, first order convergence in the timestep is observed, however, the leading error terms are, from Remark 9, of the form $O(h) + O(\Delta\tau) + O(h^2/\Delta\tau)$. As a result, convergence is ensured only if h goes to zero faster than $\sqrt{\tau}$, and for $h \sim \Delta\tau$ first order convergence is expected. Because of the maximum norm stability and linear interpolation, the error does not explode even as $\Delta\tau \rightarrow 0$ for fixed h . In fact, the solution goes to zero here as the interpolation introduces increasing artificial diffusion (see also Remark 8) and the solution is absorbed at the boundaries.
4. **Policy timestepping, linear interpolation, reference mesh:** Although not an issue for $J = 2$ control parameters, if the dimension of the switching system is J , the number of interpolations from each mesh onto all other meshes is an $O(J^2)$ operation. We can avoid this using a single ‘reference mesh’ to keep track of the solution. So in addition to the two meshes associated with V_1 and V_2 as above under item 3., we introduce a reference mesh, uniform on $[\log(K) - 4 \cdot \bar{\sigma}, \log(K) + 4 \cdot \bar{\sigma}]$, and $\tilde{V}_{k,i}$ is constructed by linear interpolation from V_k onto the reference mesh and then onto the i -th mesh point of the j -th mesh. With now two interpolations for each solution every timestep, convergence is still of first order but with a significantly higher factor than with direct interpolation between meshes. The number of interpolations needed for a J -dimensional switching system is now $O(J)$.
5. **Policy timestepping, cubic interpolation:** Finally, to reduce the accumulated interpolation error, we use the possibility of limited higher order interpolation for the mesh transfer afforded to us by Condition 1, first without the use of a reference mesh. In particular, we use monotone piecewise cubic Hermite interpolation as in [21]. Note that the interpolation [21] is monotonicity preserving but not monotone in the viscosity sense [7]. However, Condition 1 is satisfied. Note that we use $c = 0$ in this test case, which, strictly speaking, does not ensure convergence to the viscosity solution. However, it is clear from Figure 5.1 that the limiting case of $c \rightarrow 0$ does in fact converge to the viscosity solution, hence it is interesting to include the case $c \equiv 0$.

The approximation order of the interpolation method in [21] is guaranteed to be cubic only if the data are in fact monotone, and this is not the case for our initial data. Nonetheless, the error is significantly reduced compared to the linear interpolation case.

6. **Policy timestepping, cubic interpolation, reference mesh:** The results with cubic interpolation onto a reference mesh are not as accurate as for direct cubic interpolation between the computational meshes, and have a similar accuracy to the results for linear interpolation without reference mesh.

5.2 Mean-variance asset allocation

As a second example we study the mean-variance asset allocation problem as discussed in [39], following the embedding technique introduced in [29, 41]. In this example, we use the same grids for each constant policy mesh, and focus on the effects of discretization of the control. This example demonstrates that piecewise constant policy timestepping does not introduce any significant extra error compared to first order Euler timestepping with either known optimal control, or a numerical




FIGURE 5.3: The uncertain volatility test case with parameters as in Table 5.1. Shown is for different methods the \log_2 error, where $e_m = e((\Delta\tau)_m, h)$ from (5.7) is the error for timestep $(\Delta\tau)_m = 1/8 \cdot 2^{-m} \in \{1/8, \dots, 1/524288\}$ and in each plot, from top to bottom, $h = 1/4, 1/16, 1/64, 1/256, 1/1024, 1/4096$. The dotted line has slope -1 . The plots refer to, from top left lexicographically: the piecewise constant time-stepping method on a fixed mesh; the direct control method; the piecewise constant time-stepping method with: linear interpolation; linear interpolation onto a reference mesh; cubic interpolation; cubic interpolation onto a reference mesh.

optimal control obtained implicitly from the finite difference scheme. We will also see that even a fairly coarse discretization of the control admissible set yields good results. We have seen this property of discretized controls in many examples.

The method determines the pre-commitment mean variance optimal strategy [8]. Note that it is possible to develop a numerical method for solution of the time-consistent version of this problem [40]. However, since the time consistent problem can be viewed as a constrained solution of the pre-commitment problem, the time consistent solution is sub-optimal compared to the pre-commitment solution. Specifically, we consider here the sub-problem given by the equation

$$\frac{\partial V}{\partial \tau} - \inf_{p \in P} \mathcal{L}^p V = 0, \quad (5.8)$$

$$V(W, 0) = \left(W - \frac{\gamma}{2}\right)^2, \quad (5.9)$$

on $(-\infty, \infty)$, and $(0, \infty)$, with

$$L^p V = \frac{1}{2} \sigma^2 p^2 W^2 \frac{\partial^2 V}{\partial W^2} + (\pi + W(r + p\sigma\xi)) \frac{\partial V}{\partial W} \quad (5.10)$$

and either $p \in (-\infty, \infty)$ or $p \in [0, p_{\max}]$. Observe that (5.8) does not satisfy Assumption 1 since p can be unbounded. We will re-parameterize the control variable to avoid this problem.

By solving equation (5.8–5.9) for various values of the parameter γ , we can trace out the efficient frontier in the expected value, variance plane [39].

We use the standard finite difference discretization and make “maximum” use of central differences [38] whenever a positive coefficient scheme is achieved and use upwind differences only where necessary for monotonicity.

The PDE (5.8–5.9) is specified on an infinite domain. For numerical purposes, we approximate this by means of a localized problem, with approximate boundary conditions at finite values of $|W|$. We use an asymptotic approximation of V for large $|W|$. In the cases we consider, an asymptotic value for the optimal control is more directly available than for the value function. Therefore, the following solution under constant control will be applied as approximate boundary condition. More precisely, the solution to the PDE

$$\frac{\partial V}{\partial \tau} - \frac{1}{2} a^2 W^2 \frac{\partial^2 V}{\partial W^2} + (\pi + bW) \frac{\partial V}{\partial W} = 0 \quad (5.11)$$

with terminal condition (5.9) is given by

$$V(W, \tau) = \alpha(\tau)W^2 + \beta(\tau)W + \delta(\tau), \quad (5.12)$$

where $\tau = T - t$ and

$$\begin{aligned} \alpha(\tau) &= \exp((a^2 + 2b)\tau), \\ \beta(\tau) &= -(\gamma + c) \exp(b\tau) + c \exp((a^2 + 2b)\tau), \\ \delta(\tau) &= -\frac{\pi(\gamma + c)}{b} (\exp(b\tau) - 1) + \frac{\pi c}{a^2 + 2b} (\exp((a^2 + 2b)\tau) - 1) + \frac{\gamma^2}{4}, \text{ where} \\ c &= 2\pi/(a^2 + b). \end{aligned}$$

r	σ	ξ	π	W_0	T	γ	λ
0.03	0.15	0.33	0.1	1	20	14.47	1.762

TABLE 5.3: *Model parameters used in numerical experiments for mean-variance problem.*

In comparison to [39], who only derive the highest-order term, this gives an asymptotically more accurate approximation and allows us to use substantially smaller domains for the computation. Following [39], we use the parameters in Table 5.3 throughout.

We study two different cases for the permissible sets for state-variable and controls, one where $W, p \in \mathbb{R}$, and one where $W \geq 0, p \in [0, p_{max}]$.

Bankruptcy allowed, unbounded control

If bankruptcy ($W < 0$) is allowed, the PDE (5.8–5.9) holds on $(-\infty, \infty)$. In this case, a closed-form solution is known from [22], where the optimal policy is given by

$$p^*(W, t) = -\frac{\xi}{\sigma W} \left[W - \left(\frac{\gamma e^{-r(T-t)}}{2} - \frac{\pi}{r} (1 - e^{-r(T-t)}) \right) \right]. \quad (5.13)$$

Moreover, under this optimal policy, we find from the formulae in [22],

$$\begin{aligned} \text{Var}[W_T] &= \frac{e^{-\xi^2 T}}{1 - e^{-\xi^2 T}} \left[E[W_T] - \left(W_0 e^{rT} + \frac{\pi(e^{rT} - 1)}{r} \right) \right]^2, \\ E[W_T] &= \left(W_0 + \frac{\pi}{r} \right) e^{-(\xi^2 - r)T} + \frac{\gamma(1 - e^{-\xi^2 T})}{2} - \frac{\pi}{r} e^{-\xi^2 T}, \end{aligned}$$

such that $(\sqrt{\text{Var}[W_T]}, E[W_T]) = (0.794, 6.784)$, and $E[(W_T - \gamma/2)^2] = \text{Var}[W_T] + E[W_T]^2 - \gamma E[W_T] + \gamma^2/4 = 0.8338$ for the parameters in Table 5.3.

As the optimal policy in the form (5.13) is unbounded, we perform the control discretization in a different control variable. Noting that p^*W is bounded as $W \rightarrow 0$, it seems natural to consider pW as control variable in this area; however, $p^*W \sim -\xi/\sigma W$ as $|W| \rightarrow \infty$. This leads us to consider

$$q = \frac{pW}{\max(1, \omega|W|)} \quad (5.14)$$

as control variable for some $\omega > 0$, and

$$\tilde{\mathcal{L}}^q V = \frac{1}{2} \sigma^2 q^2 \max(1, \omega^2 W^2) \frac{\partial^2 V}{\partial W^2} + (\pi + Wr + q \max(1, \omega|W|) \sigma \xi) \frac{\partial V}{\partial W}. \quad (5.15)$$

The optimal control $q^*(W, t)$ will be bounded on a localized domain, and we fix an interval $Q = [q_{min}, q_{max}]$ in which we search for the optimal control by a crude approximation. In this whole process, a precise knowledge of the exact optimal control is not necessary, as we only use the rough asymptotic shape. For the computations below, we pick $\omega = 5$ and $Q = [-2.5, 3.5]$. Since we solve the PDE on a localized domain ($|W|$ bounded), and the control is now bounded as well, the localized version of equations (5.8–5.9) now satisfies Assumption 1.




FIGURE 5.4: *The mean-variance test case with parameters as in Table 5.3. Left: The numerical approximation to $V(W, 0)$ for piecewise constant policies for $M = 80$, $N = 80$ and 20 policy steps. The dashed line is the asymptotic approximation for $|W| \rightarrow \infty$. Right: The corresponding approximation to the optimal policy $q(W, 0)$, and the analytical optimal policy (5.13). The horizontal lines are at the asymptotic optimal policies for $|W| \rightarrow \infty$.*

We note that this is an example where the optimal control is an unbounded function of the state variable, but by a suitable reformulation the piecewise constant policy timestepping method can still be applied, with the policy chosen from a bounded set.

From (5.13) one sees that $q^*(W, t) \rightarrow -\xi/\sigma$ for $|W| \rightarrow \infty$. Therefore, asymptotically, (5.8), (5.10) takes the form (5.11) with suitable a and b , obtained by inserting the constant asymptotic optimal policy. We can then use the asymptotically exact boundary conditions (5.12) for both W_{max} and W_{min} . We choose $W_{max} = 40$ and $W_{min} = -40$ in the computations.

The discretized switching system has the form

$$\frac{V_j^{n+1} - \min_{1 \leq k \leq J} (V_k^n - c_{k,j})}{\Delta \tau} - L_{q_j}^h V_j^{n+1} = 0, \quad j = 1, \dots, J, \quad (5.16)$$

where $c_{k,j}$ is defined as in equation (5.6). In this case, we can set the switching parameter $c_{k,j} = 0$ since no interpolation is used, and this reduces to conventional piecewise constant policy timestepping [27]. Then the numerical approximations to all J components of the switching system are the same in each timestep after the minimum is taken.

Fig. 5.4 shows the value function V and its asymptotic approximation for large $|W|$. The two functions have visually identical tangents at the boundaries, and indeed experimentation with the values of W_{min} and W_{max} shows that the results around $W = 1$ are not significantly affected by this approximation.

Also shown in Fig. 5.4 is the approximate optimal policy obtained numerically from the policy timestepping discretization with 20 policy steps, and the exact formula (5.13), transformed into a bounded control as per (5.14).

Table 5.4 illustrates the convergence as the control mesh is refined for a fixed time and spatial mesh. The estimated order of convergence over these refinement levels is 2. We pick $J = 40$ fixed for the following tests of the convergence in the mesh size and timestep. For this value, the control

	$J = 5$	$J = 8$	$J = 10$	$J = 15$	$J = 20$	$J = 29$	$J = 40$	$J = 57$	$J = 80$
(a)	2.257	1.531	1.429	1.254	1.230	1.196	1.186	1.180	1.178
(b)		-0.725	-0.101	-0.175	-0.0241	-0.0339	-0.0104	$-5.4 \cdot 10^{-3}$	$-2.5 \cdot 10^{-3}$
(c)			7.12	0.58	7.26	0.71	3.25	1.92	2.16

TABLE 5.4: The mean-variance test case with parameters as in Table 5.3. Convergence of the control discretization alone for $N = 480$, $M = 120$. Shown are: (a) the numerical solution $V_k = \hat{V}(N, M, J_k; W_0, 0)$; (b) the increments $V_k - V_{k-1}$; (c) the ratios $(V_k - V_{k-1})/(V_{k-1} - V_{k-2})$, for $J_k = \lceil 5 \cdot \sqrt{2}^{k-1} \rceil$.




FIGURE 5.5: The mean-variance test case with parameters as in Table 5.3. Similar to Fig. 5.3 in the previous section, the \log_2 error for $(\Delta\tau)_m = 1/8 \cdot 2^{-m} \in \{1/8, \dots, 1/32768\}$ and in each plot, from top to bottom, $h = 1/4, 1/16, 1/64, 1/256, 1/1024$. The straight line has slope -1 . Left: Piecewise constant policy timestepping with 40 equally spaced policies in $[-2.5, 3.5]$. Right: Using the exact policy given by (5.13).

discretization error was empirically negligible (compared to the time and spatial discretization error).

Fig. 5.5 shows the convergence of the approximations for piecewise constant policy timestepping and for the use of the exact policy given by (5.13). In the latter case, the error is solely due to the Euler time-discretization and spatial finite differences. For piecewise constant timestepping, 40 policies were used, so that the computational time for the same mesh is about a factor of 40 larger than for a single policy, hence we show slightly fewer refinements. It appears that the spatial approximation error for large mesh size h is smaller if knowledge of the optimal control is used. The envelope showing the time discretization error is consistent with first order convergence. Interestingly, the intercept is about 4 units higher for the exact policy, so that the results using policy timestepping are about a factor 16 more accurate for the same timestep. This is not to be expected generally and must result from opposite signs of the Euler truncation error and the error due to piecewise constant policies.




FIGURE 5.6: *The mean-variance test case with parameters as in Table 5.3, with no bankruptcy and bounded control. All model parameters are as in Table 5.3. Left: Value function $V(W, 0)$ (solid line) and the asymptotic approximation (dashed line) for large W . Right: The numerical optimal policy $p(W, 0)$.*

No bankruptcy, bounded control

If bankruptcy ($W < 0$) is not allowed, the PDE (5.8–5.9) holds on $(0, \infty)$. The boundary equation at $W = 0$ is then

$$V_\tau(0, \tau) - \pi V_W(0, \tau) = 0, \quad (5.17)$$

see [39] for a discussion. For $\pi > 0$ there is an outgoing characteristic (going backwards in time) so that no boundary condition is required and we can approximate (5.17) by upwind differences from interior mesh points. In fact, as we are switching to upwind differences locally whenever the monotonicity of the scheme is violated (see above and [38]), upwinding will always be used for small W if $\pi > 0$.

For bounded control with no short-selling, $P = [0, p_{max}]$ in (5.8). In the computations, we choose $p_{max} = 1.5$ as an attained upper bound (the same used in [39]). This would correspond to a typical leverage constraint. For large W , we use again the approximation (5.12), with coefficients based on the asymptotic optimal control $p = 0$ (see Fig. 5.6).

The numerically computed value function (a closed-form solution is not available in this case) is shown in Fig. 5.6, together with the asymptotic approximation for large W . Also shown is the numerically computed optimal control.

We compare the results achieved by piecewise policy timestepping to those achieved by the direct control formulation. For clarity, the two discretizations used are

$$\frac{V^{n+1} - V^n}{\Delta\tau} - \min_{q \in Q_h} L_q^h V^{n+1} = 0 \quad (5.18)$$

for the direct control method and (5.16) for the piecewise constant timestepping method.

We use policy iteration as in [9] to solve the discrete control problem in (5.18). We use a positive coefficient discretization [38] with central differencing used as much as possible. For the direct control method, and the piecewise constant policy timestepping method, it is straightforward

		$M = 800$ $N = 50$ $J = 5$	$M = 1600$ $N = 100$ $J = 8$	$M = 3200$ $N = 200$ $J = 10$	$M = 6400$ $N = 400$ $J = 15$	$M = 12800$ $N = 800$ $J = 20$
Policy timestepping	(a)	1.5930	1.5589	1.5447	1.5378	1.5350
	(b)		-0.0341	-0.0141	-0.0069	-0.0028
	(c)			2.41	2.04	2.45
Direct control	(a)	1.5902	1.5577	1.5442	1.5376	★
	(b)		-0.0326	-0.0135	-0.0066	★
	(c)			2.4167	2.0333	★
Fixed control ($q = 1.5$)	(a)	3.4268	3.4199	3.4140	3.4104	3.4085
	(b)		-0.0069	-0.0059	-0.0036	-0.0019
	(c)			1.1733	1.6523	1.8399

TABLE 5.5: *The mean-variance test case with parameters as in Table 5.3, with no bankruptcy and bounded control. Shown are, for the policy timestepping method, the direct control method, and for a fixed constant control: (a) the numerical solution $V_k = \tilde{V}(N_k, M_k, J_k; W_0, 0)$; (b) the increments $V_k - V_{k-1}$; (c) the ratios $(V_k - V_{k-1})/(V_{k-1} - V_{k-2})$ for $M_k = 800 \cdot 2^{k-1}$, $N_k = 50 \cdot 2^{k-1}$, $J_k = \lceil 5 \cdot \sqrt{2^{k-1}} \rceil$ (except for the fixed control case, where $J = 1$).*

to verify that the discretization is monotone, consistent and stable [38, 27]. The results are shown in Table 5.5.

In each step of the policy iteration, the maximum (over parameters) of the discretized differential operator at any given mesh-point has to be computed. As the discretization (local upwinding based on the coefficients) depends on the control parameter in a discontinuous way, this maximum is found by discretizing the control and exhaustive search. This makes the complexity of a single policy iteration identical to a single timestep of the constant policy timestepping algorithm. Thus, overall, the typically observed 4–6 iterations in every timestep translates into a 4–6 factor of increase in the CPU cost of the direct control method compared to the piecewise constant policy timestepping technique. Due to this increased cost, we do not show the direct control results for the finest level (marked ★).

The refinements were chosen such that at the coarsest level a single separate refinement of the spatial mesh, timestep and control mesh gave comparable (empirical) accuracy improvements. This ensures that the data test the convergence order in all three discretization parameters. It is clear that the achieved accuracy is almost identical for both methods.

We also include results for the value achieved with a fixed control, $q = 1.5$. This is the chosen upper bound and the optimal value attained in an interval around $W_0 = 1$, see Fig. 5.6. The results are distinctly different from those under the optimal control, which shows that the similar performance of policy timestepping and direct control is not a result of the control being constant near $W = 1$. The errors for fixed control are purely due to the time and spatial finite difference discretization, and are slightly smaller than those observed in the true optimal control problems.

6 Conclusions

This article analyzes the piecewise constant policy timestepping method both from a theoretical and an applications perspective. Our main result is that if we use different meshes for each constant

policy PDE solve, then convergence to the viscosity solution can be proven even if high order (not necessarily monotone) interpolation techniques are used. Essentially, this is because we can view the piecewise constant policy timestepping method as the solution to a switching system of PDEs, where the coupling between the PDEs occurs only in the zeroth order term. However, this generality comes at a price: we must include a finite switching cost in the switching system. Convergence to the solution of the original HJB PDE occurs only in the limit as the switching cost tends to zero. However, our numerical experiments show that good results are obtained for very small (even zero) switching costs.

The general approach we follow also has superficial similarities with the “semi-Lagrangian methods” (SLM) of [14] and [18]. They both make use of the fact that for *given* coefficients (controls), it may be easier to construct monotone schemes together with the underlying mesh, especially in more than one dimension. If different controls require different meshes, interpolation of the mesh solution is needed in every timestep. In the present method this serves to carry out the optimization over solutions with different policies.

The computational results demonstrate that a smaller error is obtained using the high order interpolation, compared to linear interpolation.

In many practical situations, the local optimization problem at each node is determined by discretizing the control and using exhaustive search. In this case, our tests show that piecewise constant policy timestepping is more efficient than standard direct control methods, as a similar level of accuracy is achieved with less computational effort. This is simply due to the fact that piecewise constant policy timestepping is unconditionally stable, and does not require a policy iteration to solve nonlinear discretized equations.

The use of piecewise constant policy timestepping can be useful in situations where generic monotone schemes are hard to construct, e.g., in multidimensional settings, whose implementation we do not consider here and leave for future work.

Finally, we note that it is straightforward to implement piecewise constant policy timestepping in existing linear PDE solution software. Hence these existing algorithms can be easily converted to solve nonlinear HJB equations.

A Proof of Lemma 1

We provide here a proof of Lemma 1,

Proof. By insertion one gets

$$\begin{aligned} |F(\mathbf{x}, \phi(\mathbf{x}), D\phi(\mathbf{x}), D^2\phi(\mathbf{x})) - F_h(\mathbf{x}, \phi(\mathbf{x}) + \xi, D\phi(\mathbf{x}), D^2\phi(\mathbf{x}))| &= \left| \sup_{q \in Q_h} L_q(\phi + \xi) - \sup_{q \in Q} L_q\phi \right| \\ &\leq \left| \sup_{q \in Q_h} L_q(\phi + \xi) - \sup_{q \in Q} L_q(\phi + \xi) \right| + \left| \sup_{q \in Q} L_q(\phi + \xi) - \sup_{q \in Q} L_q\phi \right| \end{aligned}$$

by the triangle inequality. From Assumption 1 and the compactness of Q , then the supremum of $L_q\phi$ is attained, say at q^* , and then

$$L_{q^*}\phi - r_{q^*}\xi = L_{q^*}(\phi + \xi) \leq \sup_{q \in Q} L_q(\phi + \xi) \leq \sup_{q \in Q} L_q\phi + \sup_{q \in Q} L_q\xi = L_{q^*}\phi + \sup_{q \in Q} (-r_q)\xi,$$

hence

$$\left| \sup_{q \in Q} L_q(\phi + \xi) - \sup_{q \in Q} L_q \phi \right| \leq \xi \sup_{q \in Q} |r_q|. \quad (\text{A.1})$$

Now let q_ξ^* be the maximizer of $L_q(\phi + \xi)$. We also have by (2.1) that there is $q_h^* \in Q_h$ with $|q_h^* - q_\xi^*| \leq h$. By uniform continuity of the coefficients in q on the compact set Q , there exists a function $\bar{\omega}_1$ so that

$$\|\sigma_{q_\xi^*} \sigma_{q_\xi^*}^T - \sigma_{q_h^*} \sigma_{q_h^*}^T\| + \|\mu_{q_\xi^*} - \mu_{q_h^*}\| + |r_{q_\xi^*} - r_{q_h^*}| + |f_{q_\xi^*} - f_{q_h^*}| \leq \bar{\omega}_1(\mathbf{x}, h) \rightarrow 0, h \rightarrow 0$$

with the usual vector and matrix norms, and hence

$$\left| L_{q_\xi^*}(\phi + \xi) - L_{q_h^*}(\phi + \xi) \right| \leq \bar{\omega}_1(\mathbf{x}, h) \max(1, |\xi|, |\phi|, \|D\phi\|, \|D^2\phi\|) \equiv \max(1, |\xi|) \bar{\omega}_2(\mathbf{x}, h)$$

for a suitably defined $\bar{\omega}_2$. Using also that $Q_h \subset Q$,

$$\sup_{q \in Q} L_q(\phi + \xi) - \max(1, |\xi|) \bar{\omega}_2(\mathbf{x}, h) \leq \sup_{q \in Q_h} L_q(\phi + \xi) \leq \sup_{q \in Q} L_q(\phi + \xi),$$

so that

$$\left| \sup_{q \in Q_h} L_q(\phi + \xi) - \sup_{q \in Q} L_q(\phi + \xi) \right| \leq \max(1, |\xi|) \bar{\omega}_2(\mathbf{x}, h) \leq \bar{\omega}_2(\mathbf{x}, h) + \frac{\bar{\omega}_2^2(\mathbf{x}, h)}{2} + \frac{\xi^2}{2}. \quad (\text{A.2})$$

From Assumption 1, and noting that the equation coefficients are uniformly continuous in q , it is easily shown that the right hand side of (A.2) is locally Lipschitz in \mathbf{x} , independent of h . The result then follows from (A.1) and (A.2), with an appropriate choice of ω_h and ω_ξ . \square

References

- [1] R. Almgren and N. Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3:5–40, 2001.
- [2] M. Avellaneda, A. Levy, and A. Parás. Pricing and hedging derivative securities in markets with uncertain volatilities. *Applied Mathematical Finance*, 2:73–88, 1995.
- [3] G. Barles. Solutions de viscosité et équations elliptiques du deuxième ordre. Lecture notes, University of Tours, 1997.
- [4] G. Barles and E.R. Jakobsen. On the convergence rate of approximation schemes for Hamilton-Jacobi-Bellman equations. *ESAIM:M2AN*, 36(1):33–54, 2002.
- [5] G. Barles and E.R. Jakobsen. Error bounds for monotone approximation schemes for Hamilton-Jacobi-Bellman equations. *SIAM Journal on Numerical Analysis*, 43(2):540–558, 2005.
- [6] G. Barles and E.R. Jakobsen. Error bounds for monotone approximation schemes for parabolic Hamilton-Jacobi-Bellman equations. *Mathematics of Computation*, 76(240):1861–1893, 2007.

- [7] G. Barles and P.E. Souganidis. Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic Analysis*, 4(3):271–283, 1991.
- [8] S. Basak and G. Chabakauri. Dynamic mean-variance asset allocation. *Review of Financial Studies*, 23:2970–3016, 2010.
- [9] O. Bokanowski, S. Maroso, and H. Zidani. Some convergence results for Howard’s algorithm. *SIAM Journal on Numerical Analysis*, 47(4):3001–3026, 2009.
- [10] M. Boulbrachene and M. Haiour. The finite element approximation of Hamilton-Jacobi-Bellman equations. *Computers & Mathematics with Applications*, 41(7):993–1007, 2001.
- [11] A. Briani, F. Camilli, and H. Zidani. Approximation schemes for monotone systems of non-linear second order differential equations: Convergence result and error estimate. *Differential Equations and Applications*, 4:297–317, 2012.
- [12] C. Burgard and M. Kjaer. Partial differential equation representations of derivatives with bilateral counterparty risks and funding costs. *The Journal of Credit Risk*, 7:Fall:75–93, 2011.
- [13] C. Burgard and M. Kjaer. Funding strategies, funding costs. *Risk*, pages 82–87, December 2013.
- [14] F. Camilli and M. Falcone. An approximation scheme for the optimal control of diffusion processes. *Modélisation mathématique et analyse numérique*, 29(1):97–122, 1995.
- [15] R. Carmona, editor. *Indifference Pricing*. Princeton University Press, Princeton, 2009.
- [16] M.G. Crandall, H. Ishii, and P.-L. Lions. User’s guide to viscosity solutions of second order partial differential equations. *Bulletin of the AMS*, 27(1):1–67, 1992.
- [17] M.A. Davis and A.R. Norman. Portfolio selection with transaction costs. *Mathematics of Operations Research*, 15(4):676–713, 1990.
- [18] K. Debrabant and E.R. Jakobsen. Semi-Lagrangian schemes for linear and fully non-linear diffusion equations. *Mathematics of Computation*, 82(283):1433–1462, 2013.
- [19] W.H. Fleming and H.M. Soner. *Controlled Markov processes and viscosity solutions*, volume 25. Springer Science & Business Media, 2006.
- [20] P.A. Forsyth and G. Labahn. Numerical methods for controlled Hamilton-Jacobi-Bellman PDEs in finance. *Journal of Computational Finance*, 11(2):1–44, 2007/2008.
- [21] F.N. Fritsch and R.E. Carlson. Monotone piecewise cubic interpolation. *SIAM Journal on Numerical Analysis*, 17:238–246, 1980.
- [22] B. Højgaard and E. Vigna. *Mean-variance portfolio selection and efficient frontier for defined contribution pension schemes*. Research Report Series. Department of Mathematical Sciences, Aalborg University, 2007.
- [23] H. Ishii and S. Koike. Viscosity solutions for monotone systems of second order elliptic PDEs. *Comm. Partial Differential Equations*, 16:1095–1128, 1991.

- [24] H. Ishii and S. Koike. Viscosity solutions of a system of nonlinear elliptic PDEs arising in switching games. *Funkcialaj Ekvacioj*, 34:143–155, 1991.
- [25] I. Karatzas. On the pricing of American options. *Applied Mathematics and Optimization*, 17(1):37–60, 1988.
- [26] N.V. Krylov. On the rate of convergence of finite-difference approximations for Bellman’s equations. *Algebra and Analysis, St. Petersburg Mathematical Journal*, 9(3):245–256, 1997.
- [27] N.V. Krylov. Approximating value functions for controlled degenerate diffusion processes by using piece-wise constant policies. *Electronic Journal of Probability*, 4(2):1–19, 1999.
- [28] N.V. Krylov. On the rate of convergence of finite difference approximations for Bellman’s equations with variable coefficients. *Probability Theory and Related Fields*, 117:1–16, 2000.
- [29] D. Li and W.-L. Ng. Optimal dynamic portfolio selection: multiperiod mean variance formulation. *Mathematical Finance*, 10:387–406, 2000.
- [30] T. Lyons. Uncertain volatility and the risk-free synthesis of derivatives. *Applied Mathematical Finance*, 2(2):117–133, 1995.
- [31] F. Mercurio. Bergman, Piterbarg and beyond: pricing derivatives under collateralization and differential rates. Working paper, Bloomberg, 2013.
- [32] R.C. Merton. Lifetime portfolio selection under uncertainty: the continuous time case. *Review of Economics and Statistics*, 51(3):247–257, 1969.
- [33] D.M. Pooley. *Numerical Methods for Nonlinear Equations in Option Pricing*. PhD thesis, University of Waterloo, 2003.
- [34] D.M. Pooley, P.A. Forsyth, and K.R. Vetzal. Numerical convergence properties of option pricing PDEs with uncertain volatility. *IMA Journal of Numerical Analysis*, 23:241–267, 2003.
- [35] R.C. Seydel. Impulse control for jump-diffusions: viscosity solutions of quasi-variational inequalities and applications in bank risk management. PhD Thesis, Leipzig University, 2009.
- [36] I. Smears and E. Süli. Discontinuous Galerkin finite element approximation of Hamilton–Jacobi–Bellman equations with Cordes coefficients. *SIAM Journal on Numerical Analysis*, 52(2):993–1016, 2014.
- [37] J. Van Der Wal. Discounted Markov games: generalized policy iteration. *Optimization Theory and Applications*, 25:125–138, 1978.
- [38] J. Wang and P.A. Forsyth. Maximal use of central differencing for Hamilton-Jacobi-Bellman PDEs in finance. *SIAM Journal on Numerical Analysis*, 46:1580–1601, 2008.
- [39] J. Wang and P.A. Forsyth. Numerical solution of the Hamilton-Jacobi-Bellman formulation for continuous time mean variance asset allocation. *Journal of Economic Dynamics and Control*, 34:207–230, 2010.
- [40] J. Wang and P.A. Forsyth. Continuous time mean variance asset allocation: a time consistent strategy. *European Journal of Operational Research*, 209:184–201, 2011.

- [41] X. Zhou and D. Li. Continuous time mean variance portfolio selection: A stochastic LQ framework. *Applied Mathematics and Optimization*, 42:19–33, 2000.